Generative AI and Data Protection

Hannah Ruschemeier

forthcoming in: Carlo/Poncibo/Ebers/Zou, Handbook for Generative AI and the Law, Cambridge University Press, 2025.

Abstract

Generative AI has catapulted into the legal debate through the popular applications ChatGPT, Bard, Dall-E and others. While the predominant focus has hitherto centred on issues of copyright infringement and regulatory strategies, particularly within the ambit of the AI Act, it is imperative to acknowledge that generative AI also engenders substantial tension with data protection laws. The example of generative AI puts a finger on the sore spot of the contentious relationship between data protection law and machine learning built on the unresolved conflict between the protection of individuals, rooted in fundamental data protection rights and the massive amounts of data required for machine learning, which renders data processing nearly universal. In the case of LLMs, which scrape nearly the whole internet, this training inevitably relies on and possibly even creates personal data under the GDPR. This tension manifests across multiple dimensions, encompassing data subjects' rights, the foundational principles of data protection, and the fundamental categories of data protection. Drawing on ongoing investigations by data protection authorities in Europe, this paper undertakes a comprehensive analysis of the intricate interplay between generative AI and data protection within the European legal framework.

Keywords: Generative AI, Data Protection, GDPR, LLMs, Regulation, Privacy

1. Prelude: risks and challenges of generative AI

Now that the initial hype around generative AI in the form of pooled language models and image generators has subsided, legal issues are coming to the fore. In addition to discussions about generative AI and copyright, there is an increasing focus on the friction between generative models and the requirements of data protection law. In the US, several lawsuits are underway against Google and OpenAI regarding potential privacy violations by generative models.[1] Regulators are currently active in the European Union[2] and the European Data Protection Board (EDPD) has set up a task force to deal with ChatGPT (EDPB, 2023). The Italian data protection authority Garante had already opened a case against OpenAI in 2023, which led to a temporary national ban on ChatGPT (GPDP, 2023). After the proceedings were concluded, the authority found violations of the GDPR (GPDP, 2024). Investigations into data protection violations are also underway in Poland (UODO, 2023). Other countries such as Germany have issued requests for information (LDI NRW, 2023),[3] and the French data protection authority has developed an action plan (CNIL, 2023). In the case of Maximilian Schrems, the data protection NGO NYOB filed a complaint with the Austrian data protection authority in April 2024 (noyb, 2024). The core issue was that ChatGPT provided incorrect information about an individual, which OpenAI neither corrected nor responded to the request for information about what information was processed. These cases make it clear that data protection authorities are already AI regulators and generative AI is a core issue for data protection.[4]

---

[1] See the class action against OpenAI: https://storage.courtlistener.com/recap/gov.uscourts.cand.425482/gov.uscourts.cand.425482.1.0.pdf.

[2] For an overview of ongoing investigations of data protection authorities outside the European Union, see Zanfir-Fortuna (2023).

[3] The State Commissioner for Data Protection and Freedom of Information Rhineland-Palatinate, LfDI asks about ChatGPT, 24 April 2023. The State Commissioner for Data Protection Schleswig-Holstein sent out a list of questions: https://www.datenschutzzentrum.de/uploads/chatgpt/20230419_Request-OpenAI_ULD-Schleswig-Holstein_IZG.pdf.

[4] Until the beginning of 2024, OpenAI did not have an establishment in Europe, so no member state data protection authority was responsible under Art. 56, 60 GDPR. Therefore, the supervisory authorities of the member states

From a legal perspective, generative models introduce a range of distinct issues, which are well-documented across various scholarly sources (Charlotin, 2023; Chen, 2024; Chiara, 2023; Gillotte, 2020; Hazell, 2023; Marcus & Pullin, 2023; El-Mhamdi et al., 2023; Sag, 2023; Weidinger et al., 2022). Particularly the foundation models on which the popular Large Language Models (LLMs) are built pose new security risks and vulnerabilities that need to be addressed. As a result, a socio-technical assessment is needed to understand these risks and the necessary safety mechanisms, including legal and ethical aspects. The understanding of risks posed by LLMs requires a contextual approach: normative rules, like law, always operate in context.

A major concern is the protection of personal data and privacy. Different experiments have shown that it is possible to extract personal and sensitive information about individuals from LLMs (Carlini et al., 2023; El-Mhamdi et al., 2023; Gupta, 2023; Nasr et al., 2023). Researchers have proven that LLMs are able to memorise training data, either through over-application of abundant parameters to small datasets, thus reducing the capacity to generalize to new data, or through optimization for generalization in long-tailed data distributions (Novelli et al., 2024). Although this phenomenon most often occurs where duplicates exist in the training data, it will still appear where training data has been partially deduplicated. Larger models with more parameters 'remember' more data than smaller models (Carlini et al., 2022). People's privacy and right to data protection are equally violated by incorrect information as by correct information they do not want published (Mühlhoff & Ruschemeier, 2024a). These risks are exacerbated by unregulated and therefore uncontrolled secondary downstream use of the models. In the case of popular LLMs operated by global technology companies, commercial resale seems remote, as the companies have no interest in giving up their exclusive economic utilization option. The situation is different for smaller, but in some cases no less risky models: Mixtral 8x7B competes with and surpasses GPT 3.5 in some respects, due to a smart architecture that combines eight different expert models, and has been recently made open source (Hacker, 2023). It is therefore all the more necessary to obtain an overview of the purposes for which these models are used and to categorize them to enable a context-based risk assessment (Mühlhoff & Ruschemeier, 2024b).

Data protection law gives rise to its own particular frictions, from the general function and technical specificities of big data applications and generative AI on the one hand, and on the other, the particularities of generative models. Generative models are used in different contexts for different purposes, to generate text, code, video, images, audio, etc. In this chapter, I will focus on LLMs, which generate text by calculating the probability of word order. Data that is linguistically translatable, i.e., can be understood by human recipients, can undoubtedly constitute personal data in the sense of data protection law. For this reason, they are a good example of the problems of how data protection law works in relation to AI generated content.

This chapter is structured as follows: first, I outline the overarching lines of conflict between data protection law and generative AI. I then go into the specific legal issues of the GDPR: The scope and legal basis for authorisation of different steps of data processing by generative AI (2), the principles of data processing (3), the rights of data subjects (4) and questions of responsibility (5). In section 6, I discuss the transferability of the argument to models that create images, audio and video. The chapter concludes with an outlook (7).

---

could all act in accordance with Art. 55 GDPR within their areas of competence. OpenAI now operates an office in Dublin, which is designated as the data controller.

2. Structural Challenges of Generative AI for Data Protection Law

Data protection law in the EU is primarily characterised by the GDPR.[5] The current system of the GDPR is rooted in the Data Protection Directive from 1995,[6] the right to data protection (article 8 CFR), the right to privacy (article 7 CFR)[7] and the primary legal foundations in article 16 TFEU.[8] The subject matter and scope of the GDPR is the processing of personal data with the objective of protecting fundamental rights and freedoms of natural persons, article 1 GDPR. Correspondingly, the understanding of 'processing' personal data is very broad and takes an all-encompassing approach to cover practically any interaction with personal data.[9] For this reason, when personal data is involved, all stages in the lifecycle of an AI model may fall within the scope of the GDPR.

From a regulatory perspective, the various steps of data processing are therefore important in the lifecycle of an AI model. With regard to generative models, the following steps can be differentiated from a legal perspective:[10] The first step is the collection of many data points as training data. These data points may comprise personal or non-personal information. In certain instances, this process utilises extremely large datasets, making it challenging, if not impossible, to differentiate between various categories of data. For instance, ChatGPT was developed using copious amounts of data freely available on the internet. The second step is the actual training of the model using the collected data, resulting in a configured model. The third step is model application, meaning that the trained model is applied to specific cases or individuals, making the model a tool that computes a specific output in response to input data (Mühlhoff & Ruschemeier, 2024b). This breath of data and the training process mean model output contains information about cases or individuals as well as that of 'third parties' that were not part of the training data.

2.1 Quantity

The first problem area relates to how the training of powerful AI models, or the processing of large amounts of data, works in relation to the amount of data processed (Zarsky, 2017). From this arises the general (unsolved) problem of AI and Data Protection, which is the sheer quantity of data processed by potent AI models.[11] Generative AI models are typically trained on billions, if not hundreds of billions, of parameters and require large amounts of training data and computing power (Brown et al., 2020). Data protection law on the other hand is based on the idea that the individual steps of data processing and the data processed can be identified. This concept applies the idea of individual control to empower individuals by allowing them to manage their own personal information (Solove, 2024). But models trained on unprecedentedly large data sets make it impossible to manually identify or even review legal compliance of the processed data and therefore harbour potential for privacy and data protection violations.[12] Furthermore, this approach conflicts with the principle of data minimisation laid down in article 5(1)(c). This mode of operation reveals the the problems of governance arising from the systematic design of the GDPR", which, for example, envisages individual consent as the basis for

---

[5] Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation) (Text with EEA relevance) OJ 2016 No. L119, 4 May 2016.

[6] Directive 95/46/EC of the European Parliament and of the Council of 24 October 1995 on the protection of individuals with regard to the processing of personal data and on the free movement of such data OJ L 281, 23.11.1995, p. 31–50.

[7] I understand privacy as one of the protected interests under data protection law. More about the relationship: Eichenhofer (2021, p. 47 et seq.).

[8] Additionally, article 39 TEU lays down specific provisions for the area of the Common Foreign and Security Policy.

[9] See V. P. & P. D. Hert, 'GDPR Art. 4 Abs. 2 Processing' (2023), par. 3.

[10] These steps are developed in Mühlhoff & Ruschemeier (2024a, 2024b).

[11] Generative models are based on neural networks, usually Generative Adversarial Networks and Transformer Networks (Huang et al., 2024).

[12] On the relationship between privacy and data protection: Gellert (2023); Kokott & Sobotta (2013).

authorisation and presupposes the identification of individual data subjects and the data to be attributed to them.

2.2 Purposes

Privacy and data protection also seem to be at odds with the general concept of generative AI when it comes to the relevance of purposes. Data protection is highly contextual, and its level of protection depends on the type of data processed, by whom, in which settings, and for which purposes, article 5(1)(b). LLMs on the other hand, cover a wide range of purposes, applications, and operating environments. According to article 3(63) of the new regulation on artificial intelligence[13] (in the following: AI Act), a general-purpose AI model includes AI models trained with a large amount of data using self-supervision at scale, which display significant generality, are capable of competently performing a wide range of distinct tasks regardless of how the model is placed on the market, and that can be integrated into a variety of downstream systems or applications. It does not include AI models used for research, development, or prototyping activities before they are placed on the market. This definition is a good description of the current market situation; OpenAI, for example, now offers a wider variety of different GPTs for specific tasks: the laundry buddy for laundry specific questions about stains and laundry settings, the sous chef that provides users with recipes, or the negotiator that helps a user to argue in their favour. These downstream applications will gain more relevance, as it can be expected that the foundation models will not be used primarily as isolated applications, as has been the case to date, but will be integrated into other models as modular building blocks. This will increase both desirable and undesirable effects due to the possible scaling of model output. This design aspect of LLMs is difficult to reconcile with the GDPR's purpose limitation principle. In particular, when models are made available to numerous third parties via an interface, ensuring the purposes for which that model, and therefore its data, are compatible with the original purposes of data collection (article 6(4)) becomes difficult, if not impossible (Mühlhoff & Ruschemeier, 2024b).

3. Scope of the GDPR and Legal Basis

To be effective, the material and geographical scope of the GDPR needs to be open (article 3.1), which can be affirmed for all generative models in use in the Union. This is followed by the question of the legal basis for the various data processing steps (article 3.2).

3.1. Scope of Application

3.1.1. Personal Data

The GDPR applies to the processing of personal data (article 2(1)) if none of the exceptions in paragraphs 2–3 apply. The specific data processing steps covered include the collection of training data and the training of the models, even before putting the system to use or placing it on the market to generate output based on user requests.

Based on the general understanding of prevalent LLMs, it is evident that personal data is processed during step one with the collection of training data, as the effectiveness of these models increases with the breadth and diversity of their data bases. In practice, it is highly likely that personal data will be incorporated (Ruschemeier, 2023a), as training datasets are created by scraping a large number of websites and extracting large amounts of text from essays, books and website entries. As a result, this information includes, for example, names, dates of birth, and other personal content that can be associated with individual people. Alongside obviously personal information such as real names, personal data can also exist when individuals can be identified through additional information.[14] This processing is covered by the GDPR, even before the actual release of the model.

---

[13] P9_TA(2024)0138 (COM(2021)0206 – C9-0146/2021 – 2021/0106(COD)).

[14] In its judgment on the Breyer case, (CJEU C-582/14, ECLI:EU:C:2016:779) the ECJ also considers whether the identification process is lawful in the context of the question of the identifiability of persons, i.e., whether personal data are involved. This, however, contradicts the protective direction of the GDPR, which aims to protect data subjects from unlawful processing of their data. For this: Hacker (2021).

The second step of data processing, the training of the model, is trickier when it comes to the identification of personal data, as a trained model may differ from its training data. An artificial neural network, for example, is represented by a large matrix of numbers, which in turn are determined by weights and other parameters such as activation thresholds (Mühlhoff & Ruschemeier, 2024b, p. 27; Abadi et al., 2016). The data in the model is not always personal data, even though training data is clearly personal data. Personal data can be anonymised model data if, for example, techniques such as differential privacy and federated machine learning are used during the training process (Abadi et al., 2016; Dwork, 2006; El Emam et al., 2015). These state-of-the-art anonymisation techniques make it possible to remove references to the training data. Where such anonymisation makes the reconstruction of training data impossible or highly unlikely, the model itself does not constitute personal data. As the current popular large language models may intentionally or unintentionally produce persistent information, it must be assumed this was not the case: research into this 'remembering' phenomenon is ongoing. From the perspective of the GDPR, different data processing steps are relevant here: the storage of the model also constitutes data processing under the GDPR if the model data is not anonymised. In addition, many authors assume that the anonymisation of personal data is also a processing operation that requires justification under the GDPR (El Emam & Alvarez, 2014; Schreurs et al., 2008; Roßnagel, 2021).[15]

Processing step number three, the production of output, can produce personal data. For this, it doesn't matter if the information provided is correct or not. So, when LLMs produce outputs that contain the names and bibliographical information of real people, they are processing personal data. Additionally, individuals can often be easily identified from the context of the text prompt or text output, or by using search engines. LLMs linked to search engines may also facilitate identification. Particularly in the case of public LLMs, it is likely that many data subjects can be identified for the reasons mentioned above (Pesch & Böhme, 2023). It is important to note that the people in the training data are not theoretically the same as the people produced in the output data even where they have the same name, as LLMs can also generate names of existing people, for example, information which can then be assigned by users.

3.1.2. Territorial scope of application

Article 3(1) of the GDPR states that the regulation applies 'to the processing of personal data in the context of the activities of an establishment of a controller or a processor in the Union, regardless of whether the processing takes place in the Union or not'. Thus, the processing of personal data does not have to take place in the Union itself, but can be performed on servers that are for example based in the USA and other third-countries. Furthering the *lex loci solutiones* principle of the GDPR, article 3(2) GDPR provides that the requirements apply even if the data processor is not located in the EU, but offers their services to EU citizens. This therefore brings global technologies such as LLMs and other AI models such as Chat-GPT, Bard, Gemini, etc, squarely under the GDPR where these are accessible from the European Union.

3.2 Legal Basis for Data Processing

All processing of personal data within the scope of the GDPR requires a legal basis, article 6(1) GDPR. The question of the legal basis for data processing across the life cycle of a generative AI system poses different problems, as it depends on the stage of the data processing. As argued before, it is essential to distinguish between the different steps of data processing when analysing AI and data protection (Mühlhoff & Ruschemeier, 2024a).

3.2.1. Collection of Training Data

The first step in the life cycle of a generative model is the collection of training data. In the case of LLMs like GPT-4 or Bard, this step consists of scraping data from the Internet. Since the big language models have scoured almost the entire Internet, the legal basis for consent, article 6(1)(a), is logically excluded. The absence of legal obligations or contractual relationships between the operators of LLMs

[15] Article 29 Working Party, Opinion 05/2014 on Anonymisation Techniques (WP 216) 0829/14/EN. On this dispute in the context of predictive analytics: Mühlhoff & Ruschemeier (2024a), p. 19.

and all internet users worldwide, leaves only the legal basis of legitimate interest for the scraping of training data, article 6(1)(f) GDPR (Zarsky, 2017; Borgesius et al., 2017).

Article 6(1)(f) GDPR states that data processing is lawful if it is necessary for the purposes of pursuing the legitimate interests of the data controller or by a third party, provided these interests do not override the fundamental rights and freedoms of the data subjects requiring protection (Ruschemeier, 2023a). According to the ECJ, that provision lays down three cumulative conditions for the processing of personal data covered by that provision to be lawful: first, the pursuit of a legitimate interest by the controller or by a third party; second, the need to process personal data for the purposes pursuing that legitimate interest; and, third, that the interests or fundamental freedoms and rights of the data subject do not take precedence over the legitimate interest of the controller or of a third party.[16]

The fact that this constitutes the only plausible legal basis already reveals the structural problem of data protection law and data-intensive technologies. Whether article 6(1)(f) provides a sufficient legal basis is determined on a case-by-case basis (Donnelly & McDonagh, 2019; Novelli et al., 2024). There are indications that general interest may outweigh the purpose of processing, or that this can be assumed if data subjects could reasonably expect processing for training purposes. Ultimately, it depends on the individual case. However, the way mass data scraping works makes it almost impossible to identify individual interests at all, and therefore cannot provide satisfactory answers in terms of legal doctrine and legal systems.

Legitimate Interest

The term ‘legitimate interests‘ is deliberately broad, encompassing a wide range of interests, whether they are legal, economic, or idealistic, excluding only hypothetical and public interests. The collection of data for training a generative model for commercial use is initially a legitimate economic interest and protected by the freedom to conduct a business under article 16 ECFR. The argument that the ECJ also cited freedom of information (article 11 (2) ECFR) as a legitimate interest in processing transferable to generative model training (Krönke, 2023) in the Google Spain case[17] does not apply to models that are only accessible for a fee. Furthermore, search engines and generative models are not comparable due to their mode of operation. Source references in search engines can be deleted or corrected, whereas LLMs generate a new text for each question for which a new probability is calculated. If the output text is incorrect, it cannot be corrected for future outputs.

Necessity

The necessity test under this provision means that the processing of personal data must be a proportionate means of achieving the legitimate interests. Processing is necessary if personal data must be processed in order to achieve the objective of the processing: in this case, a trained AI system. Exceptionally, processing may not be necessary if only anonymised data are sufficient to train the model; this is not the case for generative models, even if the data could be anonymised in the training phase under certain circumstances.

Balancing of Interests

Article 6(1)(f) GDPR requires a balancing of the conflicting rights and interests, which must also consider the rights of data subjects under articles 7 and 8 ECFR. Their interests are particularly affected when AI collects, links, and contextualises personal data available on the Internet in response to user queries.

The fact that powerful generative models require a large amount of training data to achieve a certain level of performance, e.g. to generate word sequences that correspond to human language, speaks in favour of the interests of the operators. However, for the general development of generative models, it

[16] Case C-252/21 Meta ECLI:EU:C:2023:537, par. 106; Case C-597/19, EU:C:2021:492, para. 106.
[17] Case C-131/12 Google Spain, ECLI:EU:C:2014:317.

is not absolutely necessary to scrape data on a scale that covers almost all publicly accessible resources on the Internet. Data sets can be generated in other ways, such as through data donations, effective consent solutions, or data collection by the data controller itself, although none of these alternative options would be able to create such a broad data set. The question is therefore which specific interest of the operator is worth protecting. Meta, for example, has publicly admitted that the acquisition of licences for copyrighted material would have made the development of generative models considerably more difficult, simply because it would have made them more expensive. The same applies to the collection of training data in a privacy-compliant manner: this approach would have required considerable resources. However, if an interest is based on structural infringements, its protective value is considerably reduced. In the Meta case[18], the ECJ also ruled that the personalisation of content—Meta's core business model—was not necessary for the operation of a social network.

The ECJ decided that legitimate interests do not adequately justify Meta's practices of tracking and profiling individuals for the purpose of conducting its behavioural advertising business across its social platforms: 'it is important to note that, despite the fact that the services of an online social network such as Facebook are free of charge, the user of that network cannot reasonably expect that the operator of the social network will process that user's personal data, without his or her consent, for the purposes of personalised advertising. In those circumstances, it must be held that the interests and fundamental rights of such a user override the interest of that operator in such personalised advertising by which it finances its activity, with the result that the processing by that operator for such purposes cannot fall within the scope of point (f) of the first subparagraph of Article 6(1) of the GDPR'.[19] This raises substantial doubts about the ability of companies like OpenAI to defend the processing of vast amounts of personal data to establish a commercial generative AI enterprise, particularly given that such tools pose numerous emerging risks to identified individuals, including issues like disinformation, defamation, identity theft, and fraud, among others.

Context is therefore critical in safeguarding privacy and data protection (Ruschemeier, 2023a). While data might be publicly accessible on the Internet, perhaps even disclosed by the data subjects themselves, this does not completely negate their legitimate interest in its protection. As recital 47 states, the interests and fundamental rights of the data subject may in particular override the interest of the data controller where personal data are processed in circumstances, or in ways that data subjects do not reasonably expect. Although it is known that data are processed in other ways on the Internet, it is also a question of the specific purposes of the processing. For example, the legitimate expectation of privacy does not mean that decades-old or deleted posts, personal websites and entries can always be used to train commercial models. It is reasonable to suggest that the typical internet user does not expect, nor intend, for their data to be utilized as training material for LLMs. Therefore, the use of the data for training these models represents a secondary purpose. In most instances, it is unlikely that publicly available data was intended to serve as a dataset for the financial gain of LLM providers. Consequently, the use of such publicly available data infringes upon contextual privacy (Nissenbaum, 2011).

Moreover, the determination of what constitutes a legitimate interest cannot be made in isolation but must consider the broader legal and regulatory context. The broad scope of scraping also means that an unmanageable number of people are affected. According to the German constitutional doctrine of the Federal Constitutional Court, this can intensify interventions if a particularly large number of people are affected without cause. This impact is referred to as 'scattering width', and is a line of argumentation also used by the ECJ.[20] Such effects also arise in the case of universal data processing, as almost all internet users are affected.

[18] Case C-252/21 Meta ECLI:EU:C:2023:537, para. 102 et seq.
[19] Case C-252/21 Meta ECLI:EU:C:2023:537 para.117.
[20] Case C-511/18 u.a., ECLI:EU:C:2020:791, Rn. 143 f. – La Quadrature du Net; Case C-293/12, ECLI:EU:C:2014:238, para. 57 ff. – Digital Rights Ireland Ltd; Case C-203/15, ECLI:EU:C:2016:970, para 105 f. – Tele2; Case C-140/20, ECLI:EU:C:2022:258, Rn. 66 – Commissioner

Additionally, the legitimate interest must also be lawful, meaning it should conform to all applicable laws and regulations, including the principles and other provisions of data protection law. This includes ensuring that processing aligns with the expectations of the data subject based on their relationship with the controller, adheres to the principles of data minimisation, and implements appropriate safeguards. In the case of broad web scraping, individual interests are difficult to identify. However, there were concerns about the legality of scraping from the outset, including with regard to potential copyright infringements. An interest pursued through structural infringements cannot be legitimate.

Compatibility with the principles of data protection law under article 5 GDPR also plays a role in the balancing of interests. These raises concerns in relation to the fairness of processing article 5(1)(a), purpose limitation (article 5(1)(b)), data minimisation (article 5 (1)(c)) and data accuracy (article 5(d)).

Therefore, it is not convincing to assume legitimate interests per se for the entire breadth of the training data. It becomes particularly complicated when data being processed pertains to minors or special categories of personal data according to article 9(1) and are included in the training data. Neither can be completely excluded. To complicate this matter further, it has not been conclusively clarified when the processing of personal data 'reveals' special categories of personal data under article 9(1) GDPR.

3.2.3 Training of the Model

In terms of training the model, a chronological review of the various data processing operations is recommended. If the data is anonymised in the course of training the model, the prevailing view is that an authorisation basis is required. In this context, anonymisation should be understood as normative and not technical. This is because the ECJ considers anonymisation has occurred even if it is technically possible, but unlikely, that the controller could carry out identification with the means available.[21] In addition, according to the court, data is considered anonymous under the GDPR if re-identification is illegal.[22]

In principle, the anonymisation of personal data is generally easy to justify under article 6 GDPR. As an expression of the principle of data minimisation and storage limitation, effective and permanent anonymisation can be used to pursue the interests of both data subjects and data controllers: the former are protected from unauthorised interference with their fundamental data protection rights, while the latter are freed from the often perceived rigid requirements of data protection law (Hornung & Wagner, 2020). This argument falls in the face of the volume of data, as effective consent from the data subjects pursuant to articles 6(1)(a) and 7 GDPR, cannot be obtained in practice.

Legal obligations to anonymise training data are conceivable, but not yet relevant in practice, article 6(1)(c). This means that the legal basis of legitimate interest from article 6(1)(f) also remains for anonymisation. As a rule, this provision can provide adequate results: anonymisation is permissible given that in the typical case it is in the interest of the data subjects themselves, or at least they will not have a conflicting interest. In the case of large LLMs, there will be no individual interest of the data subject in non-anonymisation, and moreover, any interest the data subject may have will not be outweighed by the other relevant interests, e.g., of the other data subjects.

Assessing the processing of special categories of personal data under article 9 is more difficult. If anonymisation is considered to be processing requiring justification, a case under article 9(2) would have to be made. As described above, the processing of special categories of data cannot be ruled out for LLMs. The hurdles of article 9(2) are high, but the 'made public' provision of article 9(2)(e) can also be considered here. Others argue for a teleological reduction of article 9(1) for anonymisation (Hornung & Wagner, 2020). In the case of anonymisation of training data, no data subject's interests are infringed

---

of An Garda Síochána.

[21] Case C-582/14, ECLI:EU:C:2016:779 – Breyer, para. 45–49.

[22] Case C-582/14, ECLI:EU:C:2016:779 – Breyer, para. 46. Rightfully critical: Hacker (2021).

under either variant. These complex considerations alone show that the GDPR, due to its individual-based approach, is not an effective tool for regulating generative AI.

3.2.3 Generating Output

Subsequently, the output of generative language models can represent the processing of personal data. A distinction must be made here between the processing of scraped training data and the processing of user data, the latter are the prompts entered while using the model. There is no legitimate interest in processing user data, e.g., in the context of input prompts when using LLMs. Instead, effective consent, pursuant to article 6(1)(a) must be obtained, a tool that needs to be critically evaluated in the digital space (Ben-Shahar & Schneider, 2011; Borgesius et al., 2017; De & Imine, 2020; Nguyen et al., 2022; Mühlhoff & Ruschemeier, 2024b, p. 31). Open AI had to update its privacy policy for EU users after being investigated by the Italian data protection authority. It now states: 'We use the content you provide to improve our services, such as to train the models that run our services. Read our instructions on how to opt out of the use of your content to train our models.' (OpenAI, 2023). However, consent is only a valid basis for prompts containing personal information about the user themselves. If users generate prompts that include personal data from other persons, they cannot validly consent on their behalf (Novelli et al., 2024).

Generated output, including personal data, means the problem of training data persists at the end of a generative model's life cycle. If there was no legal basis for collecting the training data, there is no legal basis for using it to generate output. Theoretically, legitimate interest could also be considered here under article 6(1)(f), assessed on a case-by-case basis according to the criteria described above. The problem is that LLMs do not allow a case-by-case assessment because of the quantity of data. In addition, generative models are scalable in terms of their output, which means that false information can be disseminated to a large number of users and third parties.

Output processing is also problematic in cases where models infer or disclose special categories of personal data under article 9(1). It has been shown that models can memorise and reproduce private and personal information such as phone numbers, addresses, and medical documents (Carlini et al., 2023). In the age of big data, it is now potentially possible to infer sensitive information from almost any data, especially if one includes the boundless category of political opinions, included in article. 9(1) GDPR. This means that 'normal' personal data can reveal special categories of personal data covered by article 9(1), although the criteria for distinguishing between general and sensitive data remain contested (Mühlhoff & Ruschemeier, 2024a, p. 22). One proposed criterion goes to the intention behind the data processing. In scenarios involving context-specific information, the purpose of the evaluation could lead to the generation of sensitive data (Matejek & Mäusezahl, 2019; Schulz, 2018; Georgieva & Kuner, 2020, p. 373; Wachter & Mittelstadt, 2019).[23] This has some support in the courts: the ECJ seemed to apply a broad interpretation of 'revealing' in the Meta case,[24] and in another ruling, the Court decided that the disclosure of a spouse, partner, or cohabitee's name could potentially indicate the sexual orientation of the applicant.[25] The court has established minimal criteria for 'revealing' sensitive data: the act of an 'intellectual operation involving comparison or deduction' is deemed sufficient to extend the special protection regime meant for sensitive data to personal data that are not inherently sensitive. However, this judgment was not directly related big data, leaving the distinction somewhat ambiguous (Mühlhoff & Ruschemeier, 2024a, p. 24).

Consequently, in many instances involving big data, the mere potential to infer sensitive information may subject processes such as AI training to the provisions of article 9, with little likelihood that LLMs satisfy the exceptions in article 9(2). The research exemption under article 9(2)(j), for instance, is

[23] Question No. 2 in ECJ Case C- 252/21 Meta Platforms and Others [2021] (ECLI:EU:C:2022:704), dismissed by the Advocate general in his Opinion, Par. 41.
[24] Case C-252/21, Meta, ECLI:EU:C:2023:537, para. 73.
[25] Case C-184/20 ECLI:EU:C:2022:601.

restricted to the development of models for research purposes and does not permit their commercial exploitation, as indicated in Recitals 159 and 162 (Novelli et al., 2024).

When applied to large language models, it depends on whether the output can be used to infer sensitive information about individuals that they have not made public themselves. Even if certain indicators, e.g., of political orientation, are available on the Internet, LLM output may aggregate this information. As such, article 9(2)(e) does not constitute a legal basis for this type of derivation.[26]

In terms of output, there are requirements for data accuracy, article 5(1)(d) GDPR. Language models have been shown to 'hallucinate' and produce incorrect information, including incorrect personal data (Wachter et al., 2024). Under the GDPR, operators are responsible for ensuring data accuracy, articles 5(2), 24, 25(1) GDPR.[27] Although there are disclaimers in all popular applications that the models may not always be correct, whether these notes have an effect remains an open question given the *automation bias* (Ruschemeier, 2023b; Hondrich & Ruschemeier, 2023; Ruschemeier & Hondrich, 2024). Even if the current error rate of LLMs (Metz, 2023) does not justify a general prohibition of such applications on the basis of ensuring data accuracy, it does affect data subjects' rights. If the right to rectification or erasure cannot be effectively enforced, data inaccuracy becomes all the more important.

4. Data Subject's Rights

As with other areas of data-intensive technology application, there are problems with the enforcement of data subjects' rights in the case of generative models (Solove, 2023). In general, especially regarding the actors involved, many data driven AI technologies are developed, promoted, sold, and used by a handful of big-tech-companies, which establishes an informational power asymmetry between the powerful processors and the users. As a result, privacy rights alone are insufficient to address the issue of data disempowerment.[28] Individuals typically possess limited capacity to manage their personal data, as there is a fundamental limit to the control they can exert. While rights can afford a modest degree of influence in certain isolated cases, this influence is too sporadic and disjointed to significantly safeguard privacy. Ultimately, rights function primarily as a minor element within a broader framework (Solove, 2023).

The sheer quantity of the data processed from various sources seems to make it impossible to identify and inform individuals of the processing, or of the processor, to enable data subjects to assert their rights with regard to the processing of their data This therefore realistically rules out compliance with the data subject's right to information. In practice, it was reported that companies such as OpenAI and Midjourney did not respond to requests for information from people who had found themselves in the training data (Harlan & Brunner, 2023).

The prerequisite for the exercise of the data subjects' rights under the GDPR provided in articles 13–22 is, first and foremost, that the data subject is aware of the data processing. Users providing input into an AI modern in the form of prompts are covered by article 13 GDPR. However, where data has not been collected from the data subjects themselves, article 14 GDPR comes into play as well. According to both standards, data subjects must be informed about who has processed which data (categories of data), for what purposes, on what legal basis, and whether this data has been disclosed to third parties. The purpose of these transparency provisions is to enable data subjects to exercise their other rights, such as the right to erasure or rectification. Article 14(5) states the exception that the transparency obligation does not apply where and insofar the provision of such information proves impossible or would involve a 'disproportionate effort, in particular for processing for archiving purposes in the public interest,

[26] See also Case C-252/21, Meta, ECLI:EU:C:2023:537, para. 73.

[27] As the ECJ already stated in 2008 with regard to the maintenance of data in the Central Register of Foreigners - similar to an actor acting under public authority, in order to delete or correct incorrect information without delay. Case C-524/06, 2008 I-09705. Krönke (2023).

[28] We have explained in detail elsewhere that systemic, supra-individual solutions and a collective understanding of data protection are needed: Mühlhoff & Ruschemeier (2024c); Mühlhoff & Ruschemeier (2024a); Mühlhoff & Ruschemeier (2024b); Mühlhoff (2023).

scientific or historical research purposes or statistical purposes […] In such cases the controller shall take appropriate measures to protect the data subject's rights and freedoms and legitimate interests, including making the information publicly available'.

Again, it depends on the individual case, although it seems doubtful whether LLM operators can invoke unreasonableness if they already knew before the model was developed that individual requests for information could not be enforced. In any case, the principle of responsibility in article 5(2) GDPR means a failure to respond to such requests or a reference to general impossibility is not sufficient.

The example shows that LLMs and possibly other generative AI models that produce content operate almost universally, not just at an individual level. This near-universal infringement reflects the profound mismatch between data-intensive models and the individual rights approach to data protection taken by data protection laws. As a result of this universality, other rights of data subjects such as the right to rectification (article 16 GDPR) and the right to erasure (article 17 GDPR) exist on paper, but become unenforceable in practice (Novelli et al., 2024).[29] Furthermore, removal requests from an individual data subject may be insufficient, particularly in cases where the same information has been disseminated by multiple users interacting with the LLM (Novelli et al., 2024; Brown et al., 2020). In essence, simply deleting data from a training dataset offers a superficial remedy, as it does not guarantee the elimination of the ability to retrieve that data or extract related information embedded within the model's parameters (Novelli et al., 2024). As the output of certain machine learning models is shaped by the data used during the training phase, the original training data or information related to removed data may be deduced or 'leaked' (Novelli et al., 2024).

6. Responsibility

In addition to the varying steps of data processing in generative models, multiple parties could potentially be considered as data controllers under the GDPR due to their various levels of involvement. The GDPR establishes three categories of responsibility for data processing in relation to the data subject: controller, processor, and third parties.

The data controller is primarily responsible for compliance with the provisions of the GDPR, article 5(2). This can also be seen in the definition of article 4(7) GDPR according to which the data controller is a 'natural or legal person [...], which, alone or jointly with others, determines the purposes and means of the processing of personal data.' Article 4(8) goes on to define the processor as a 'natural or legal person [...] to which the personal data are disclosed, whether a third party or not.' Third parties on the other hand are actors other than the data subject, controller or processor, article 4(10).

Prima facie, the legal companies that develop and deploy generative models are the data controllers. However, a differentiated picture emerges in the various steps of data processing. Indisputably, companies like OpenAI and Google act as data controllers in relation to the processing steps involved in establishing the training foundation and storing model parameters, given they exclusively determine the modalities of data processing, such as the decision to release a freely accessible LLM. However, in terms of output production, generative models process data based on the prompts from their users. Whether this can be used to establish joint responsibility of providers and users within the meaning of article 26 GDPR remains an open question.

Joint responsibility under article 26 GDPR refers to the situation where two or more controllers jointly determine the purposes and means of data processing. In contrast, the relationship of a controller and a processor is different (article 28 GDPR); in this constellation, processing is carried out *on behalf* of the controller. Joint responsibility is thus a relationship of equality, whereas the data processor is subject to

[29] Novelli et al. (2024) rightfully point out, that invoking the right to erasure (and correction under article 21 GDPR) depends on whether the LLM itself is personal data. This depends whether the training data is anonymised or not, see 3.2.3.

the instructions of the data controller. Whether these evaluations can be transferred to the relationship between providers of generative models and users is questionable.

Users are not considered processors within the meaning of article 28 GDPR, as while there is a contract between them and the providers they do not have the obligations of a processor, especially those of article 28(3) GDPR. However, the purpose of generative models is to enable users to freely use the model for their own defined purposes, free from instructions from the providers, e.g., to formulate letters, find cooking recipes, or revise texts.

Users and providers could therefore be joint controllers, requiring the purposes of data processing be jointly defined and mutual obligations set out transparently. The classification of joint controllership is supported by the fact that users and providers both influence the purposes of data processing: the providers of generative models provide the basic framework for the purposes of use, which are specified by users. Both therefore have a reciprocal effect on data processing and are also dependent on each other. This is however contradicted by the fact that users also tend to be data subjects. Article 26(3) GDPR provides that data subjects can act against each of the joint controllers. Although the parties involved do not have to hold the same level of responsibility, mere contributory causation without cooperative action is not sufficient for joint responsibility.[30] Additionally, users of generative models cannot effectively be held responsible to third parties because they do not influence data processing to the same extent. They have no ability to grant rights of access to providers or to delete personal data from the training data.

The relationship between users and providers of generative AI is therefore a special case that cannot be seamlessly subsumed under the categories of the GDPR. On the one hand, users are more than just data subjects, as their active inputs are required to generate and shape the model's output. On the other hand, they are neither data processors nor joint controllers, as they have no influence over the fundamental modalities of data processing. For instance, providers are able to simply deactivate models or make them subject to fees (as in the case of ChatGPT). The ECJ considers the extent to which data controllers participate in the joint data processing and within which concrete processing phases to be crucial.[31] In this case, users only participate in output generation, which is significantly dependent on the previous steps, such as training. The purpose and aim of the regulations concerning joint control is to counteract a diffusion of responsibility among multiple participants. Affected individuals should be able to clearly identify who is collecting their personal data, and for what purpose (recital 58). Therefore, the providers may be responsible for user-generated content in the case of generative models, but the inverse does not apply. This follows from the reasoning and fundamental rights protection of the GDPR provisions regarding responsibility and also corresponds to the technical and economic reality.[32]

5. Images, Audio and Videos as Personal Data

The considerations for LLMs are not always transferable to generative models that produce audio, images, and video. This is because the aim of these models is not to generate information, which may be incorrect, but to generate new audio or visual material. The primary aim, of generating new content, has given rise to many copyright issues arising in these cases. Images and videos can also be personal data if they can be used to identify the person, something easily achieved today, through image searches.

A major problem is the significant increase in deepfakes in the digital context, which now affects not only public figures but also the general population. Women in particular are often victims of deepfake pornography, where explicit images and videos are generated using their images, without their consent. This is an unlawful processing of personal data that violates the GDPR and, in many cases, national regulations. The photographed image of a person in these cases constitutes personal data if the person is still alive, regardless of whether the data is fake or not. The purpose of deepfakes is to disparage or

[30] Case C-2010/16 ECLI:EU:C:2018:388, para. 45.
[31] Ibid.
[32] Similiar: Case C- 131/ 12, ECLI:EU:C:2014:317 – Google Spain, para. 33.

discredit a specific individual, thus fulfilling the decisive characteristic of article 4(1) GDPR, namely that the person is identified or identifiable. Voices can be personal data, if the person is identifiable, visual or acoustic identification methods recorded using pattern recognition, such as facial or voice recognition (speaker recognition), can even be considered biometric data under article 4(14) GDPR (Weichert, 2024).

The AI Act only imposes a labelling obligation on deepfakes (article 50(4)), leaving considerable doubt as to whether there is an adequate level of legal protection at European level (Mania, 2024; Kumkar & Rapp, 2022; Fallis, 2021).

6. Conclusion and Outlook

The popular use cases of generative AI models show that data protection law is reaching its limits when it comes to regulating data-intensive technologies. In addition to the problems highlighted here, further questions arise regarding the principle of purpose limitation for data processing for data-intensive models and their downstream applications. The use of LLMs in decision-making situations raises questions about the scope of the prohibition in article 22 GDPR.

Structural problems of almost universal concern exist between the GDPR's focus on individual protection and the volume of data, and in terms of a structural enforcement deficit, particularly with regard to data protection principles and data subjects' rights.

As important as the structure of data protection law is for the protection of fundamental rights, new solutions are needed for the structural challenges posed by generative AI and other data-intensive technologies. These may also lie outside data protection law. To address these challenges, it is important to recognise the structural dimension of AI as a socio-technical development (Mühlhoff, 2020). As a result, there is a need for structural solutions that go beyond the enforcement of individual rights (Mühlhoff & Ruschemeier, 2024a, 2024b). The AI Act does not contribute solutions to the structural and specific challenges for data protection law posed by generative AI, despite its stated goal of protecting fundamental rights including data protection. This is due to its structure following product safety law which has a fundamentally different approach than legal frameworks aimed at protecting fundamental rights, such as the GDPR. The AI Act establishes certain obligations for high-risk AI-systems such as data governance, transparency requirements, and standards for human oversight. However, these provisions do not address the privacy and data protection of users. This leaves a great need for legal regulation.

Reference list

- Abadi, M., A. Chu, I. Goodfellow, H. B. McMahan, I. Mironov, K. Talwar, & L. Zhang (2016). Deep Learning with Differential Privacy: *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security - CCS*, 308–18.

- Ben-Shahar, O. & C. E. Schneider (2011). The Failure of Mandated Disclosure: *University of Pennsylvania Law Review, 159,* 647–749.

- Borgesius, F. J. Z., S. Kruikemeier, S. C. Boerman, & N. Helberger (2017). Tracking Walls, Take-It-Or-Leave-It Choices, the GDPR, and the ePrivacy Regulation: *European Data Protection Law Review, 3,* 353–68.

- Brown, T. B., B. Mann, N. Ryder, N. Ryder, M. Subbiah, J. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, P. Shyam, G. Sastry, G. Sastry, A. Askell, S. Agarwal, A. Herbert-Voss, G. Krueger, T. Henighan, R. Child, A. Ramesh, A. Ramesh, D. M. Ziegler, D. M. Ziegler, J. Wu, C. Winter, C. Hesse, M. Chen, E. Sigler, M. Litwin, S. Gray, B. Chess, J. Clark, C. Berner, S. McCandlish, A. Radford, I. Sutskever, & D. Amodei (2020). Language models are few-shot learners: *Proceedings of the 34th International Conference on Neural Information Processing Systems*, 1877–1901.

- Carlini, N., D. Ippolito, M. Jagielski, K. Lee, F. Tramer, & C. Zhang (2022). *Quantifying Memorization Across Neural Language Models*.

- Carlini, N., J. Hayes, M. Nasr, M. Jagielski, V. Sehwag, F. Tramèr, B. Balle, D. Ippolito, & E. Wallace (2023). *Extracting Training Data from Diffusion Models*. doi: 10.48550/arxiv.2301.13188.

- Charlotin, D. (2023). Large Language Models and the Future of Law: *SSRN Electronic Journal.* Retrieved from https://www.ssrn.com/abstract=4548258. doi: 10.2139/ssrn.4548258.

- Chen, S. (2024). *Potential Applications and Safety of Large Language Models in Healthcare, Interdisciplinary Humanities and Communication Studies*, *1*(6), 1–5. doi: 10.61173/f578jp05.

- Chiara, P. G. (2023). Italy: Italian DPA v. OpenAI's ChatGPT: The Reasons Behind the Investigation and the Temporary Limitation to Processing: *European Data Protection Law Review, 9*(1), 68–72. doi: 10.21552/edpl/2023/1/12.

- Commossion Nationale de l'Informatique et des Libertés (CNIL) (2023, May 16). Artificial intelligence: the action plan of the CNIL. Retrieved from https://www.cnil.fr/en/artificial-intelligence-action-plan-cnil.

- De, S. J. & A. Imine (2020). Consent for targeted advertising: the case of Facebook: *AI & SOCIETY, 35,* 1055–64.

- Donnelly, M. & M. McDonagh (2019). Health Research, Consent and the GDPR Exemption: *European Journal of Health Law, 26,* 97–119.

- Dwork, C. (2006). Differential Privacy. In M. Bugliesi, B. Preneel, V. Sassone, & I. Wegener (Eds.), *Automata, Languages and Programming*: *33rd International Colloquium, ICALP, Venice, Italy, July 10–14, 2006, Proceedings, Part II* (pp.1–12). Springer.

- Eichenhofer, J. (2021). *e-Privacy. Theorie und Dogmatik eines europäischen Privatheitsschutzes im Internet-Zeitalter: Jus Publicum*, vol. 301.

- El Emam, K. & C. Alvarez (2014). A critical appraisal of the Article 29 Working Party Opinion 05/2014 on data anonymization techniques: *International Data Privacy Law, 5,* 73–87.

- El Emam, K., S. Rodgers, & B. Malin (2015). Anonymising and sharing individual patient data: *BMJ (Clinical research ed.) 350*, h1139.

- Gillotte, J. L. (2020*). Copyright Infringement in AI-Generated Artworks*: *UC Davis Law Review, 53*(5)*, 2655–2691.*

- European Data Protection Board (EDPB) (2023, April 13): EDPB resolves dispute on transfers by Meta and creates task force on Chat GPT. Retrieved from https://www.edpb.europa.eu/news/news/2023/edpb-resolves-dispute-transfers-meta-and-creates-task-force-chat-gpt_en.

- Fallis, D. (2021). The Epistemic Threat of Deepfakes: *Philosophy & Technology, 34,* 623–43

- Garante per la protezione dei dati personali (GPDP) (2023, March 31). Intelligenza artificiale: il Garante blocca ChatGPT. Raccolta illecita di dati personali. Assenza di sistemi per la verifica dell'età dei minori. Retrieved from https://www.garanteprivacy.it:443/home/docweb/-/docweb-display/docweb/9870847.

- Garante per la protezione dei dati personali (GPDP) (2024, January 29). ChatGPT: Garante privacy, notificato a OpenAI l'atto di contestazione per le violazioni alla normativa privacy. Retrieved from https://gpdp.it/web/guest/home/docweb/-/docweb-display/docweb/9978020.

- Gellert, R. & S. Gutwirth (2013). The legal construction of privacy and data protection: *Computer Law & Security Review*, *29*, 522–30.

- Georgieva, L. & C. Kuner (2020). Art. 9 Processing of special categories of personal data. In C. Kuner, L. A. Bygrave, C. Docksey, L. Drechsler (eds.), *The EU General Data Protection Regulation (GDPR)* (pp. 365–84). Oxford University Press.

- Gupta, M., C. Akiri, K. Aryal, E. Parker, & L. Praharaj (2023). From ChatGPT to ThreatGPT: Impact of Generative AI in Cybersecurity and Privacy: *IEEE Access*, *11*, 80218–45.

- Hacker P. (2021). A legal framework for AI training data—from first principles to the Artificial Intelligence Act: *Law, Innovation and Technology 13*, 257–301.

- Hacker, P. (2023, December 13). What's Missing from the EU AI Act: Addressing the Four Key Challenges of Large Language Models: *Verfassungsblog*. Retrieved from https://verfassungsblog.de/whats-missing-from-the-eu-ai-act.

- Harlan, E. & Brunner, K. (2023, July 7). We Are All Raw Material for AI. Retrieved from https://interaktiv.br.de/ki-trainingsdaten/en.

- Hazell, Julian (2023, December 14). Spear Phishing with Large Language Models. Retrieved from https://www.governance.ai/research-paper/llms-used-spear-phishing.

- Hondrich, L. & H. Ruschemeier (2023). Addressing automation bias through verifiability: *EWAF'23: European Workshop on Algorithmic Fairness*.

- Hornung, G. & B. Wagner (2020). Anonymisierung als datenschutzrelevante Verarbeitung? Rechtliche Anforderungen und Grenzen für die Anonymisierung personenbezogener Daten: *ZD,* 223–28. Retrieved from https://www.bfdi.bund.de/SharedDocs/Downloads/DE/Konsultationsverfahren/1_Anonymisierung/Stellungnahmen/Universit%C3%A4t-Kassel.pdf?__blob=publicationFile&v=5.

- Huang, K., Y. Wang, & X. Zhang (2024). Foundations of Generative AI. In K. Huang, Y. Wang, B. Goertzel, Y. Li, S. Wright, & J. Ponnapalli (Eds.), *Generative AI Security: Theories and Practices*, 3–30. Springer.

- Kokott, J. & C. Sobotta (2013). The distinction between privacy and data protection in the jurisprudence of the CJEU and the ECtHR: *International Data Privacy Law, 3*, 222–28.

- Krönke, C. (2023, April 14). Attention Is All You Need: ChatGPT und die DSGVO: *Verfassungsblog*. Retrieved from https://verfassungsblog.de/attention-is-all-you-need.

- Kumkar, L. K. & J. P. Rapp (2022). Deepfakes: Eine Herausforderung für die Rechtsordnung: Zeitschrift für Digitalisierung und Recht (ZfDR), *3*, 199–228.

- Landesbeauftragte für Datenschutz und Informationssicherheit NRW (LDI NRW) (2023, October 30). Prüfung von ChatGPT geht in die nächste Runde. Retrieved from https://www.badische-zeitung.de/der-neue-sc-freiburg-coach-julian-schuster-kickt-beim-trainingsauftakt-direkt-selbst-mit.

- Mania, K. (2024). Legal Protection of Revenge and Deepfake Porn Victims in the European Union: Findings From a Comparative Legal Study: *Trauma, Violence, & Abuse, 25,* 117–29.

- Marcos, H. & M. Pullin (2023, October 11). Large Language Models and EU Data Protection: Mapping (Some) of the Problems — The Digital Constitutionalist (October 2023). Retrieved from https://digi-con.org/large-language-models-and-eu-data-protection-mapping-some-of-the-problems.

- Matejek, M. & S. Mäusezahl (2019). Gewöhnliche vs. sensible personenbezogene Daten. Abgrenzung und Verarbeitungsrahmen von Daten gem. Art. 9 DS-GVO: *ZD,* 551–56.

- Metz, C. (2023, November 6). Chatbots May "Hallucinate" More Often Than Many Realize. Retrieved from https://www.nytimes.com/2023/11/06/technology/chatbots-hallucination-rates.html.

- El-Mhamdi, E.-M., S. Farhadkhani, R. Guerraoui, N. Gupta, L.-N. Hoang, R. Pinot, S. Rouault, & J. Stephan (2023). *On the Impossible Safety of Large AI Models*. doi: 10.48550/arXiv.2209.15259.

- Mühlhoff, R. & H. Ruschemeier (2024a). Predictive Analytics and the Collective Dimensions of Data Protection: *Law, Innovation and Technology*, *16*(1), 261–292. doi: 10.1080/17579961.2024.2313794.

- Mühlhoff, R. & H. Ruschemeier (2024b). Regulating AI with Purpose Limitation for Models: *Journal of AI Law and Regulation, 1*, 24–39.

- Mühlhoff, R. & Ruschemeier, H. (2024c). Predictive Analysis und DSGVO: *Kommunikation & Recht (K&R), 12*, 15–26.

- Mühlhoff, R. (2020). Human-aided artificial intelligence: Or, how to run large computations in human brains? Toward a media sociology of machine learning: *New Media & Society, 22,* 1868–84.

- Mühlhoff, R. (2023). Predictive Privacy: Collective Data Protection in Times of AI and Big Data: *Big Data & Society,* 1–14.

- Nasr, M., N. Carlini, J. Hayase, M. Jagielski, A. F. Cooper, D. Ippolito, C. A. Choquette-Choo, E. Wallace, F. Tramèr, & K. Lee (2023). *Scalable Extraction of Training Data from (Production) Language Models*. doi: 10.48550/arXiv.2311.17035.

- Nguyen, T. T., M. Backes, & B. Stock (2022). Freely Given Consent? Studying Consent Notice of Third-Party Tracking and Its Violations of GDPR in Android Apps: *Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security, (New York, NY, USA: Association for Computing Machinery, 2022)*, 2369–83.

- Nissenbaum, H. (2011). A contextual approach to privacy online: *Daedalus, 140,* 32–48.

- None of your business (noyb) – Europäisches Zentrum für digitale Rechte (2024, April 29). ChatGPT provides false information about people, and OpenAI can't correct it. Retrieved from https://noyb.eu/en/chatgpt-provides-false-information-about-people-and-openai-cant-correct-it.

- Novelli, C., F. Casolari, P. Hacker, G. Spedicato, & L. Floridi (2024). Generative AI in EU Law: Liability, Privacy, Intellectual Property, and Cybersecurity. Retrieved from https://papers.ssrn.com/abstract=4694565.

- OpenAI (2023, December 15). Datenschutzerklärung. Retrieved from https://openai.com/de/policies/eu-privacy-policy.

- Pesch, P. & Böhme (2023). Verarbeitung personenbezogener Daten und Datenrichtigkeit bei großen Sprachmodellen: *Multimedia und Recht*, *16*(12), 917–23.

- Roßnagel, A. (2021). Datenlöschung und Anonymisierung. Verhältnis der beiden Datenschutzinstrumente nach der DSGVO: *ZD,* 11, 188–92.

- Ruschemeier, H. (2023a, April 7). Squaring the Circle. Retrieved from https://verfassungsblog.de/squaring-the-circle.

- Ruschemeier, H. (2023b). The Problems of the Automation Bias in the Public Sector – A Legal Perspective. In Weizenbaum Institute (ed.), *Weizenbaum Conference Proceedings: AI, Big Data, Social Media and People on the Move, (Rochester, NY, 2023b)* (pp. 59–69).

- Ruschemeier, H. & L. Hondrich (2024). Automation Bias in Public Administration - an Interdisciplinary Perspective From Law and Psychology: *Government Information Quarterly, 41*(3). doi: 10.1016/j.giq/2024.101953.

- Sag, M. (2023). Copyright Safety for Generative AI: *Houston Law Review*, *61*(2). doi: 10.2139/ssrn.4438593.

- Schulz, S. (2018). Art. 9 DS-GVO, para. 13.

- Schreurs, W. J., M. Hildebrandt, E. Kindt, & V. Michaël (2008). Cogitas, Ergo Sum. The Role of Data Protection Law and Non-discrimination Law in Group Profiling in the Private Sector. In M. Hildebrandt, S. Gutwirth (Eds.), *Profiling the European Citizen* (pp. 241–70). Springer.

- Solove, D. J. (2023). The Limitations of Privacy Rights: *Notre Dame Law Review*, *98*, 975.

- Solove, D. J. (2024). Artificial Intelligence and Privacy.

- Urząd ochrony danych osobowych (UODO) (2023, September 9). Technologia musi być zgodna z Rodo. Retrieved from https://uodo.gov.pl/pl/138/2823.

- Wachter, S. & B. Mittelstadt (2019). A Right to Reasonable Inferences: Re-thinking Data Protection Law in the Age of Big Data and AI: *Columbia Business Law Review,* 1–135.

- Wachter, S., B. Mittelstadt, & C. Russell (2024). Do large language models have a legal duty to tell the truth?

- Weidinger, L., J. Uesato, M. Rauh, C. Griffin, P.-S. Huang, J. Mellor, A. Glaese, M. Cheng, B. Balle, A. Kasirzadeh, C. Biles, S. Brown, Z. Kenton, W. Hawkins, T. Stepleton, A. Birhane, L. A. Hendricks, L. Rimell, W. Isaac, J. Haas, S. Legassick, G. Irving, & I. Gabriel (2022). Taxonomy of Risks posed by Language Models: *2022 ACM Conference on Fairness, Accountability, and Transparency*, *(Seoul Republic of Korea: ACM)* (pp. 214–29).

- Weichert, T. (2024). DS-GVO, Art. 4 Nr. 14 biometrische Daten, marg. 3.

- Zanfir-Fortuna, G. (2023, September 12). How Data Protection Authorities are De Facto Regulating Generative AI - Future of Privacy Forum. Retrieved from https://fpf.org/blog/how-data-protection-authorities-are-de-facto-regulating-generative-ai.

- Zarsky, T. Z. (2017). *Incompatible: the GDPR in the age of big data*: *Seton Hall L. Rev. 47*, 995–1020.